

23rd USENIX Conference on File and Storage Technologies (FAST'25)

Santa Clara, California, USA
25-27 February 2025

ISBN: 979-8-3313-1575-7

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2025) by Usenix Association
All rights reserved.

Printed with permission by Curran Associates, Inc. (2025)

For permission requests, please contact Usenix Association
at the address below.

Usenix Association
2560 Ninth Street, Suite 215
Berkeley, California, 94710

<https://www.usenix.org/>

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

23rd USENIX Conference on File and Storage Technologies (FAST '25)

February 25–27, 2025

Santa Clara, CA, USA

Tuesday, February 25

File Systems

Fast, Transparent Filesystem Microkernel Recovery with Ananke 1
Jing Liu, *Microsoft Research*; Yifan Dai, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau, *University of Wisconsin–Madison*

Boosting File Systems Elegantly: A Transparent NVM Write-ahead Log for Disk File Systems 19
Guoyu Wang, Xilong Che, Haoyang Wei, Shuo Chen, Puyi He, and Juncheng Hu, *Jilin University*

DJFS : Directory-Granularity Filesystem Journaling for CMM-H SSDs 35
Seung Won Yoo, *Korea Advanced Institute of Science and Technology (KAIST)*; Joontaek Oh, *University of Wisconsin–Madison*; Myeongin Cheon and Bonmoo Koo, *Korea Advanced Institute of Science and Technology (KAIST)*; Wonseob Jeong, Hyunsob Song, Hyeonho Song, and Donghun Lee, *Samsung Electronics*; Youjip Won, *Korea Advanced Institute of Science and Technology (KAIST)*

ScaleLFS: A Log-Structured File System with Scalable Garbage Collection for Commodity SSDs 53
Jin Yong Ha, *Seoul National University*; Sangjin Lee, *Chung-Ang University*; Hyeonsang Eom, *Seoul National University*; Yongseok Son, *Chung-Ang University*

Rethinking the Request-to-IO Transformation Process of File Systems for Full Utilization of High-Bandwidth SSDs 69
Yekang Zhan, Haichuan Hu, Xiangrui Yang, and Qiang Cao, *Huazhong University of Science and Technology*; Hong Jiang, *University of Texas at Arlington*; Shaohua Wang and Jie Yao, *Huazhong University of Science and Technology*

Cloud Storage

FlacIO: Flat and Collective I/O for Container Image Service 87
Yubo Liu, Hongbo Li, Mingrui Liu, Rui Jing, Jian Guo, Bo Zhang, Hanjun Guo, Yuxin Ren, and Ning Jia, *Huawei Technologies Co., Ltd.*

Cloudscape: A Study of Storage Services in Modern Cloud Architectures 103
Sambhav Satija, Chenhao Ye, Ranjitha Kosgi, Aditya Jain, Romit Kankaria, Yiwei Chen, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau, *University of Wisconsin–Madison*; Kiran Srinivasan, *NetApp*

Maat: Analyzing and Optimizing Overcharge on Blockchain Storage 123
Zheyuan He, *University of Electronic Science and Technology of China*; Zihao Li, *The Hong Kong Polytechnic University*; Ao Qiao and Jingwei Li, *University of Electronic Science and Technology of China*; Feng Luo, *The Hong Kong Polytechnic University*; Sen Yang, *University of Electronic Science and Technology of China*; Gelei Deng, *Nanyang Technological University*; Shuwei Song, Xiaosong Zhang, and Ting Chen, *University of Electronic Science and Technology of China*; Xiapu Luo, *The Hong Kong Polytechnic University*

Revisiting Network Coding for Warm Blob Storage 139
Chuang Gan, *Huazhong University of Science and Technology*; Yuchong Hu, *Huazhong University of Science and Technology and Shenzhen Huazhong University of Science and Technology Research Institute*; Leyan Zhao, Xin Zhao, Pengyu Gong, and Dan Feng, *Huazhong University of Science and Technology*

Machine Learning and Storage

MOONCAKE: Trading More Storage for Less Computation — A KVCache-centric Architecture for Serving LLM Chatbot 155
Ruoyu Qin, *Moonshot AI and Tsinghua University*; Zheming Li, Weiran He, and Jialei Cui, *Moonshot AI*; Feng Ren, Mingxing Zhang, Yongwei Wu, and Weimin Zheng, *Tsinghua University*; Xinran Xu, *Moonshot AI*

Towards High-throughput and Low-latency Billion-scale Vector Search via CPU/GPU Collaborative Filtering and Re-ranking	171
Bing Tian, Haikun Liu, and Yuhang Tang, <i>Huazhong University of Science and Technology</i> ; Shihai Xiao, <i>Huawei Technologies Co., Ltd</i> ; Zhuohui Duan, Xiaofei Liao, and Hai Jin, <i>Huazhong University of Science and Technology</i> ; Xuecang Zhang and Junhua Zhu, <i>Huawei Technologies Co., Ltd</i> ; Yu Zhang, <i>Huazhong University of Science and Technology</i>	

IMPRESS: An Importance-Informed Multi-Tier Prefix KV Storage System for Large Language Model Inference	187
Weijian Chen, Shuibing He, Haoyang Qu, Ruidong Zhang, Siling Yang, and Ping Chen, <i>Zhejiang University</i> ; Yi Zheng and Baoxing Huai, <i>Huawei Cloud</i> ; Gang Chen, <i>Zhejiang University</i>	

GPHash: An Efficient Hash Index for GPU with Byte-Granularity Persistent Memory	203
Menglei Chen, Yu Hua, Zhangyu Chen, Ming Zhang, and Gen Dong, <i>Wuhan National Laboratory for Optoelectronics, School of Computer, Huazhong University of Science and Technology</i>	

GeminiFS: A Companion File System for GPUs	221
Shi Qiu, Weinan Liu, Yifan Hu, Jianqin Yan, and Zhirong Shen, <i>NICE Lab, Xiamen University</i> ; Xin Yao, Renhai Chen, and Gong Zhang, <i>Huawei Theory Lab</i> ; Yiming Zhang, <i>NICE Lab, Xiamen University and Shanghai Jiao Tong University</i>	

Wednesday, February 26

More Machine Learning

3L-Cache: Low Overhead and Precise Learning-based Eviction Policy for Caches	237
Wenbin Zhou, <i>Beijing University of Technology</i> ; Zhixiong Niu and Yongqiang Xiong, <i>Microsoft Research</i> ; Juan Fang and Qian Wang, <i>Beijing University of Technology</i>	

LeapGNN: Accelerating Distributed GNN Training Leveraging Feature-Centric Model Migration	255
Weijian Chen, Shuibing He, and Haoyang Qu, <i>Zhejiang University</i> ; Xuechen Zhang, <i>Washington State University Vancouver</i>	

HiDPU: A DPU-Oriented Hybrid Indexing Scheme for Disaggregated Storage Systems	271
Wenbin Zhu, Zhaoyan Shen, and Qian Wei, <i>Shandong University</i> ; Renhai Chen, <i>Tianjin University and Huawei Technologies Co., Ltd</i> ; Xin Yao, <i>Huawei Technologies Co., Ltd</i> ; Dongxiao Yu, <i>Shandong University</i> ; Zili Shao, <i>The Chinese University of Hong Kong</i>	

PIMLex: A High-Performance Learned Index with Processing-in-Memory	287
Lixiao Cui, Kedi Yang, Yusen Li, Gang Wang, and Xiaoguang Liu, <i>College of Computer Science, Nankai University</i>	

Hardware Assist

HaSiS: A Hardware-assisted Single-index Store for Hybrid Transactional and Analytical Processing	305
Kecheng Huang, <i>The Chinese University of Hong Kong</i> ; Zhaoyan Shen, <i>Shandong University</i> ; Zili Shao, <i>The Chinese University of Hong Kong</i> ; Feng Chen, <i>Indiana University Bloomington</i> ; Tong Zhang, <i>Rensselaer Polytechnic Institute and ScaleFlux Inc.</i>	

AegonKV: A High Bandwidth, Low Tail Latency, and Low Storage Cost KV-Separated LSM Store with SmartSSD-based GC Offloading	321
Zhuohui Duan, Hao Feng, Haikun Liu, Xiaofei Liao, Hai Jin, and Bangyu Li, <i>Huazhong University of Science and Technology</i>	

D2FS: Device-Driven Filesystem Garbage Collection	337
Juwon Kim and Seungjae Lee, <i>Korea Advanced Institute of Science and Technology (KAIST)</i> ; Joontaek Oh, <i>University of Wisconsin–Madison</i> ; Dongkun Shin, <i>Sungkyunkwan University</i> ; Youjip Won, <i>Korea Advanced Institute of Science and Technology (KAIST)</i>	

ShiftLock: Mitigate One-sided RDMA Lock Contention via Handover	355
Jian Gao, Qing Wang, and Jiwu Shu, <i>Tsinghua University</i>	

Selective On-Device Execution of Data-Dependent Read I/Os	373
Chanyoung Park, Minu Chung, and Hyungon Moon, <i>UNIST (Ulsan National Institute of Science and Technology)</i>	

Security, Integrity, and Consistency

- On Scalable Integrity Checking for Secure Cloud Disks** 391
Quinn Burke, Ryan Sheatsley, Rachel King, Owen Hines, Michael Swift, and Patrick McDaniel, *University of Wisconsin–Madison*
- Silhouette: Leveraging Consistency Mechanisms to Detect Bugs in Persistent Memory-Based File Systems** 407
Bing Jiao, *Florida State University*; Ashvin Goel, *University of Toronto*; An-I Andy Wang, *Florida State University*
- OPIMQ: Order Preserving IO stack for Multi-Queue Block Device** 425
Jieun Kim, *Korea Advanced Institute of Science and Technology (KAIST)*; Joontaek Oh, *University of Wisconsin–Madison*; Juwon Kim, Seung Won Yoo, and Youjip Won, *Korea Advanced Institute of Science and Technology (KAIST)*
- AWUPF Rediscovered: Atomic Writes to Unleash Pivotal Fault-Tolerance in SSDs** 441
Jiyune Jeon and Jongseok Kim, *Sungkyunkwan University*; Sam H. Noh, *Virginia Tech*; Euseong Seo, *Sungkyunkwan University*
- AtomicDisk: A Secure Virtual Disk for TEEs against Eviction Attacks** 449
Hongliang Tian, *Ant Group*; Xinyi Yu, *NICE Lab, Xiamen University*; Shaowei Song and Qingsong Chen, *Ant Group*; Zhihao Zhang and Shiyu Wang, *NICE Lab, Xiamen University*; Weijie Liu, *Nankai University*; Erci Xu, *Shanghai Jiao Tong University*; Shoumeng Yan, *Ant Group*; Yiming Zhang, *NICE Lab, Xiamen University and Shanghai Jiao Tong University*

Thursday, February 27

Compression and Deduplication

- MedFS: Pursuing Low Update Overhead via Metadata-Enabled Delta Compression for Log-structured File System on Mobile Device** 461
Chao Wu and Cheng Ji, *Nanjing University of Science and Technology*; Li-Pin Chang, *National Yang Ming Chiao Tung University*; Zongwei Zhu, *University of Science and Technology of China*; Congming Gao, *Xiamen University*; Weichao Guo and Chao Yu, *Guangdong Oppo Mobile Telecommunications Corp., Ltd*; Yanzhi Wang, *Northeastern University*
- Don't Maintain Twice, It's Alright: Merged Metadata Management in Deduplication File System with GOGETAFS** 479
Yanqi Pan and Wen Xia, *Harbin Institute of Technology, Shenzhen*; Erci Xu, *Alibaba Group*; Hao Huang, Xiangyu Zou, and Shiyi Li, *Harbin Institute of Technology, Shenzhen*
- Archer: Adaptive Memory Compression with Page-Association-Rule Awareness for High-Speed Response of Mobile Devices** 497
Changlong Li, *East China Normal University*; Zongwei Zhu and Chao Wang, *University of Science and Technology of China*; Fangming Liu, *Huazhong University of Science and Technology and Peng Cheng Laboratory*; Fei Xu and Edwin H. -M. Sha, *East China Normal University*; Xuehai Zhou, *University of Science and Technology of China*
- VectorCDC: Accelerating Data Deduplication with Vector Instructions** 513
Sreeharsha Udayashankar, Abdelrahman Baba, and Samer Al-Kiswany, *University of Waterloo*

Storage Diversity and Heterogeneity

- Oasis: An Out-of-core Approximate Graph System via All-Distances Sketches** 523
Tsun-Yu Yang, *The Chinese University of Hong Kong (CUHK)*; Yi Li, *The University of Texas at Dallas*; Yizou Chen, *The Chinese University of Hong Kong (CUHK)*; Bingzhe Li, *The University of Texas at Dallas*; Ming-Chang Yang, *The Chinese University of Hong Kong (CUHK)*
- PolyStore: Exploiting Combined Capabilities of Heterogeneous Storage** 539
Yujie Ren, *Rutgers University and EPFL*; David Domingo, Jian Zhang, and Paul John, *Rutgers University*; Rekha Pitchumani, *Samsung Semiconductor Inc.*; Sanidhya Kashyap, *EPFL*; Sudarsun Kannan, *Rutgers University*
- Liquid-State Drive: A Case for DNA Block Device for Enormous Data** 557
Jiahao Zhou, Mingkai Dong, Fei Wang, Jingyao Zeng, Lei Zhao, Chunhai Fan, and Haibo Chen, *Shanghai Jiao Tong University*
- DNA data storage: A generative tool for Motif-based DNA storage** 573
Samira Brunmayr, Omer S. Sella, and Thomas Heinis, *Imperial College London*